Two Views on Time Reversal*

Jill North†‡

In a recent paper, Malament (2004) employs a time reversal transformation that differs from the standard one, without explicitly arguing for it. This is a new and important understanding of time reversal that deserves arguing for in its own right. I argue that it improves upon the standard one. Recent discussion has focused on whether velocities should undergo a time reversal operation. I address a prior question: What is the proper notion of time reversal? This is important, for it will affect our conclusion as to whether our best theories are time-reversal symmetric, and hence whether our spacetime is temporally oriented.

1. Introduction. In a recent paper, David Malament (2004) employs a time reversal transformation that differs from the standard one. Aside from noting the naturalness and general applicability of this transformation, Malament does not offer explicit argument for rejecting the standard account. This is because his purpose is to argue against David Albert's view (2000, Chapter 1) that classical electromagnetism is not time reversal invariant.

I think Malament has been too modest. He has proposed a new and important understanding of time reversal that deserves arguing for in its own right. In this paper, I argue that Malament's time reversal transformation improves upon the standard one.

Recent discussion of time reversal has focused on whether velocities should undergo a time reversal operation. I address a prior question: What is the proper notion of time reversal? The answer to this question is important, for it will affect our conclusion as to whether our best theories

*Received February 2007; revised March 2008.

†To contact the author, please write to: Department of Philosophy, Yale University, P.O. Box 208306, New Haven, CT 06520-8306; e-mail: jill.north@yale.edu.

‡For comments and discussion, I am indebted to Frank Arntzenius, Hartry Field, Stephen Leeds, David Malament, Ted Sider, students in seminars at NYU and Yale, audience members at the Pacific APA in 2006, and an anonymous referee for this journal.

Philosophy of Science 75 (April 2008) pp. 201–223. 0031-8248/2008/7502-0005$10.00
Copyright 2008 by the Philosophy of Science Association. All rights reserved.

https://doi.org/10.1086/590199 Published online by Cambridge University Press

are time reversal symmetric, and hence whether the spacetime of our world is temporally oriented.[1]

2. Why Does It Matter? What I call the standard transformation is the one that is generally assumed by physicists and philosophers. I will need to motivate my lumping together under the heading 'the standard view' such disparate positions as those of the physics books and of philosophers like Albert (2000, Chapter 1) and Paul Horwich (1987, Chapter 3), who argue against much of what the books say. I will do so in a moment.

First let us ask the obvious question: why should we care what the proper time reversal transformation is? Why not assume the one the books do, and leave it at that? Forget the details of this or that particular operation. Define our time reversal transformation, and a corresponding sense of time reversal symmetry, any way we like.

The reason we cannot be so nonchalant is that there is something important we want the transformation to do for us. We want it to tell us about the spacetime structure of our world.

In applying any transformation to a theory, we hope to learn about the symmetry of the theory, and of the world that theory describes. We do this by comparing the theory with what happens to it after the transformation. If the theory remains the same after the transformation—if it is *invariant* under the transformation—then it is symmetric under that operation. This indicates that the theory 'says the same thing' regardless of how processes are oriented with respect to the structure underlying the transformation. We conclude that a world described by the theory lacks the structure that would be needed to support an asymmetry under the operation. For example, from the space-translation invariance of the laws, we infer that space is homogeneous, that there is no preferred location in space.

As for transformations and their related symmetries generally, so too for the time reversal transformation and time reversal symmetry specifically. We are not interested in the time reversal transformation *per se*. What we are ultimately after is the nature of time itself. We want to know whether there is an objective distinction between past and future: a temporal orientation on the spacetime manifold.[2]

1. I limit discussion to classical theories. Though the general conclusions remain, the proper time reversal transformation could change when we take into account quantum field theory (Arntzenius 2004, 2005; Arntzenius and Greaves 2007).

2. A temporal orientation τ^a is an everywhere continuous timelike vector field on the (temporally orientable) spacetime manifold that contains only vectors which lie in one lobe of the light cones. τ then determines a temporal orientation on the manifold: it picks out, at each spacetime point, which direction is to the future of that point, and

https://doi.org/10.1086/590199 Published online by Cambridge University Press

We cannot directly observe whether spacetime has this structure. Nor do the phenomena suffice to tell us this. Things are asymmetrically distributed in space, but we do not conclude from this alone that space itself is asymmetric, possessing structure picking out a preferred spatial direction. Similarly, there may be asymmetries in the way things are laid out in time, but this alone does not indicate that time itself is asymmetric.[3]

We can learn about the structure of time a different way: by means of the fundamental dynamical laws. For the following inference seems reasonable: Non-time-reversal-invariant laws would give us reason to believe that spacetime is temporally oriented.

The principle behind the inference is this. If the fundamental laws cannot be formulated without reference to a particular kind of structure, then this structure must exist in order to support the laws—'support' in the sense that the laws could not be formulated without making reference to that structure. In general, we infer that spacetime has a certain structure on the basis of corresponding (non-)invariances in the fundamental laws of physics. Consider the inference to a lack of preferred spatial location. If the fundamental laws were *not* invariant under spatial translations, then we *would* infer that space is inhomogeneous. For we would need to include a preferred-point spatial structure in giving our theory of the world. The laws could not be formulated without it.

Similarly here. If the laws are *not* invariant under time reversal, then we could not state them without presupposing a temporal orientation on the spacetime manifold—an objective distinction between the two temporal directions, indicating in which one things are allowed to evolve and in which they are not. The laws themselves make reference to the distinction. If these are the fundamental laws, then we have reason to infer that the world has the structure needed to support the distinction. If the laws *are* symmetric under time reversal, then they do not presuppose a temporal orientation. They say the same thing regardless of the direction things evolve in. If these are the fundamental laws, then we would not infer a temporal orientation. In giving our theory of the world, we would not need to include this structure in its spacetime. This inference is not conclusive; but it is reasonable. (And it is generally endorsed: see, e.g., Savitt 1996; Arntzenius 1995, 1997. Still, there could be highly nonlocal laws or asymmetric boundary conditions rather than an asymmetric space-

---

which is to the past.

This is not to say the question is whether there is an objective fact as to which direction is past and which is future. The question is whether there is a structural difference between the two temporal directions, regardless of which we call or experience as 'past' and which 'future'.

3. As Price puts it (1996, Chapter 1), we are interested in the asymmetry *of* time, over and above the asymmetry of things *in* time.

https://doi.org/10.1086/590199 Published online by Cambridge University Press

time structure. The inference must ultimately consider our best fundamental theory; see note 4.)

We want to know whether there is a distinction between past and future that is independent of everything other than the structure of time itself (independent of our experience in time, of an arbitrarily chosen reference frame, and so on). The best way to learn about this is to look for relevant invariances in the fundamental laws. And the time reversal transformation is the mathematical gadget that we use to see whether the laws have these invariances.

Thus, our time reversal transformation should be able to tell us whether the laws are noninvariant in a way that suggests a direction of time. This, in turn, means that the inference to a direction of time is what should guide us in locating the proper transformation. So this is our test: the proper time reversal transformation should be such that, when we compare a theory with its time reverse, we can reasonably infer whether spacetime has a temporal orientation.

Note that we should have in hand our best, most fundamental theory when making this inference. There are temporally asymmetric, nonfundamental theories that do not suggest an underlying structural asymmetry; thermodynamics is an example. The idea is that, once we have in hand what we take to be our best fundamental theory, operating on it with the proper transformation should reveal what temporal structure is needed to support the theory, and therefore what structure we should infer the spacetime to have.[4]

4. Even then, the inference will not be conclusive. An anonymous referee offered this example. Consider two worlds, each containing a single one-dimensional object. In one world, the fundamental law says the object doubles its length per unit time. In the other, the fundamental law says the object halves its length per unit time. The law for each world is time asymmetric, so that according to our guiding inference, there is a temporal orientation in each world. But consider a different law for each of these worlds (or rather, if this is the law, then there really is no genuine distinction between the two worlds): the single object changes its length at a uniform rate per unit time. This law is time symmetric, so that according to our guiding inference, there is no temporal orientation in either world. The symmetric law seems just as much a candidate best theory for each of these worlds; both theories save the phenomena. But then how could the phenomena ever indicate whether the symmetric or asymmetric theory is correct, and therefore whether time is temporally oriented—not only in simple worlds like these, but more generally as well?

The short answer is that our guiding inference is not conclusive; but this does not mean it is unreasonable. Our conclusion about time's arrow will depend on what we take to be the best overall fundamental theory. And which theory we take to be 'best' will involve the usual cost-benefit analyses, which may themselves be inconclusive. In the example here, if we accept the symmetric theory, then we do away with some temporal structure, but at the cost of a highly nonlocal theory (the theory says something like: the object uniformly changes its length in the direction in which it had been

https://doi.org/10.1086/590199 Published online by Cambridge University Press

Summing up: What makes for a good time reversal transformation? The proper transformation should link up in the right way with the inference to a direction of time. We can indeed define other kinds of 'time reversal' transformations. But only this test will locate the one to use in figuring out whether our spacetime is temporally oriented. A good transformation is hard to find.

3. Active and Passive Transformations. There are two ways of understanding a transformation—actively or passively. It is worthwhile for what comes later to review the difference between the two.

Geometric objects (points, vectors, tensors, and the like) exist independently of any coordinate system. They take on different coordinate-dependent representations, in terms of different components, depending on the coordinate system being used. But the objects themselves are invariant under coordinate changes. Think of a vector in a two-dimensional Euclidean plane: rotate coordinate axes, and the components of the vector will change; the vector itself (think of its length given by the Pythagorean theorem) will not.

There are two kinds of transformation that we can apply to theories about these objects. An active transformation leaves the coordinate system alone and transforms the coordinate-independent objects to alter their coordinate-dependent descriptions. It is 'active' in that it acts directly on the objects to yield their different coordinate descriptions. A passive transformation instead changes the coordinates, altering objects' coordinate-dependent descriptions. It is 'passive' in that it leaves the objects alone, transforming the coordinate system around them. Actively, a solution is mapped to another solution; passively, the coordinatization is changed. Thus, actively, a vector is rotated, changing its components. Passively, the coordinate axes are rotated (in the opposite direction), changing the components in the same way.

For any such transformation, there is typically both a passive and an active version: the passive transformation yields the *same* changes to objects' coordinate-dependent representations as the corresponding active

changing all along). Or we could chalk up the observed asymmetry to asymmetric boundary conditions, which may or may not, depending on one's view, be sufficient grounds to infer a temporal orientation. (Many think not; Maudlin 2007 is a dissenting view.) If we accept the asymmetric law, on the other hand, then we get rid of the nonlocality and the positing of special boundary conditions (assuming they are not needed for other reasons) at the expense of additional spacetime structure. Which is the best theory, all things considered? It is not clear in this case. Still, there is reason to be optimistic that worlds like ours will provide enough evidential considerations to pick out one of the candidate best theories, which we can then plug into our guiding inference.

https://doi.org/10.1086/590199 Published online by Cambridge University Press

transformation does. It seems we would like this to be true of our time reversal transformation as well. If a time reversal transformation, understood actively, changes solutions to a theory into nonsolutions, then the corresponding passive transformation should do the same. After all, if the transformations did not agree in this way, that could spell trouble for our guiding inference. Why should we infer a temporal orientation structure if the passive transformation allows us to make this inference, but the active version of the transformation does not?

Our guiding inference can be put either way. Passively: non-time-reversal-invariant laws do not remain the same under a change in time coordinate from t to $-t$ (equivalently, under an inversion of the time axis about the temporal coordinate origin). Actively: take a solution to the theory—a sequence or history of states that is possible according to the laws of the theory—and apply the time reversal transformation directly to it. The theory is not time reversal invariant if doing so transforms the solution into a nonsolution (a process that is no longer possible). Either way, if this is the proper time reversal transformation, then we can reasonably infer whether spacetime is temporally oriented, according to the theory.

4. The Standard Transformation. The standard time reversal transformation is captured by the image of a movie playing backward: any film of an allowable process on a symmetric theory also depicts an allowable process when we run that film in reverse. In other words, for any sequence of states allowed by the theory, the temporally inverted sequence of states is also allowed. The running of a film backward is the playing out of the time-reversed sequence of states. Watch a movie of an allowable process, and if the theory is temporally symmetric, we would not be able to tell whether the movie is playing forward or backward.

This is an intuitive understanding of time reversal symmetry, what we initially have in mind when we think of what it means for the direction of time to matter or not to a given theory. If the time-reversed sequence of states were *not* allowed, then processes governed by the theory are sensitive to the direction of time in which they occur. We would be able to tell which way the movie is playing.

Physics texts also allow a time reversal operator to act on the instantaneous states comprising the time-reversed process. This alters the understanding of time reversal a bit. Now we say that a theory is time reversal invariant just in case, for any sequence of instantaneous states . . . , $S(t_0)$, $S(t_1)$, $S(t_2)$, . . . allowed by the theory, the reverse sequence of *time-reversed* states, . . . , $S^T(t_2)$, $S^T(t_1)$, $S^T(t_0)$, . . . , is also allowed, where T is

https://doi.org/10.1086/590199 Published online by Cambridge University Press

the appropriate time reversal operator (time runs left to right).[5] (T inverts velocities in Newtonian mechanics, for example.)[6]

Not everyone agrees with this. Albert (2000, Chapter 1) and Horwich (1987, Chapter 3) disagree with the books on what happens to objects under time reversal. They argue that the proper time reversal transformation should convert a given sequence of states into the temporally inverted sequence of the *very same states*. No time reversal operations on the instantaneous states themselves are allowed.[7]

Yet all parties agree on the basic action of the time reversal transformation as a *mirroring of states across a time*. To get the time reverse of a process, flip the states that comprise the process across the temporal origin. A theory is time reversal symmetric when this mirroring operation always transforms a possible process into another possible process.[8]

In agreement with Albert and Horwich, the textbook view considers the time-reverse of a process to be the original sequence of states (albeit acted on by T) shifted to their temporal mirror-image locations. The difference between Albert and Horwich, on the one hand, and the textbooks, on the other, lies in what kinds of objects get transformed under

5. More precisely, for any history $t \mapsto S(t)$ allowed by the theory, the time-reversed history $t \mapsto {}^{T}S(t)$ is also allowed.

6. T will vary from theory to theory. In Newtonian mechanics, T inverts velocities. Think of a given Newtonian process run backward. If velocities are included in the instantaneous states (as they standardly are, though see below), they must flip direction in order for the backward process to be possible (indeed, for it to be coherent). In classical electromagnetism, T inverts **B** fields. One reason to question the standard view is its lack of a principle on the basis of which to determine the appropriate time reversal operator, for any given theory. The worry is that the resultant notion of time reversal will be trivial. For an argument that completely free rein in defining our time reversal operators results in a trivial notion of time reversal, see Arntzenius 2004, 32–33.

7. Hence Albert's unorthodox view that classical electromagnetism is not time reversal invariant. The magnetic field would have to invert in the time-reversed process for the theory to be time reversal invariant. Yet Albert only allows time reversal operations on nonfundamental quantities, such as velocities, that are time derivatives of fundamental ones, such as particle positions in Newtonian mechanics. Horwich treats the magnetic field similarly, though he concludes that electromagnetism *is* time reversal invariant. His own view suggests he should agree with Albert on this point, however: see Arntzenius 2004, note 4. Albert and the books both say that particles and scalar field values get shifted to their temporal mirror image locations, and that the directions of the velocities get flipped. But their justification for this differs. For Albert, since velocities are defined as time derivatives of positions, they will invert by virtue of taking the time-reversed sequence of positions, which for him are fundamental quantities in classical theories.

8. This criterion must be modified a bit for indeterministic theories. See Arntzenius 1997 for discussion and a clarification of the standardly cited criterion.

https://doi.org/10.1086/590199 Published online by Cambridge University Press

this action—not in the basic mirroring action of the transformation itself. Henceforth, I take the 'standard view' to be the transformation that mirrors objects across a time, and consider this to encompass the position of the textbooks and of Albert and Horwich. This will allow us to focus on the primary innovation of Malament's proposal.

Note that the standard transformation, as just stated, is an active one: it shifts objects an equal temporal distance across the time origin. Though the books focus on this understanding of the transformation, they note a corresponding passive version. This is simply a change in time coordinate from t to $-t$, or an inversion of the time axis about the temporal origin. This will yield the same changes to objects' coordinate-dependent descriptions that the standard active transformation does. (At least for objects in classical space; in Section 7 we will see a problem for this idea applied to the objects in spacetime.) On both versions of the transformation, particles and field values get shifted an equal temporal distance across the time origin, for example: actively, the objects themselves get shifted that way; passively, they get the correspondingly shifted time coordinate.

5. The Standard Transformation, Revised and Updated. Malament approaches the issue of time reversal by means of the invariant formulation of a given theory. This is an illuminating way of going about things. I wish to follow suit.

One might wonder why we should take this approach; discussions in physics books generally do not. The reason is simple. Our theories of physics are about coordinate-independent objects, things like particles and fields. These objects can take on different descriptions depending on the coordinate system we use, but they are physical entities that exist in the world independently of how we choose to describe them. What is more, our best physical theories tell us that different choices of coordinates can yield equally legitimate ways of describing these objects. This suggests that the real, physical features of the world must be the ones that hold regardless of the coordinate system being used: the ones that hold in all allowable coordinate systems or reference frames.[9]

Hence our best theories are typically formulated in an invariant, coordinate-independent way. This is the way best suited to getting at the nature of the world, apart from our descriptions of it. Mathematically, these theories are formulated in terms of invariant geometric objects, rather than coordinate-dependent numerical ones—in terms of vectors, for example, rather than their components. And when we consider the

9. I have been influenced here by Maudlin (1994, Chapter 2; 2006).

https://doi.org/10.1086/590199 Published online by Cambridge University Press

time reversal (or other) symmetry of a theory, we should consider our best, most fundamental theory, in its best formulation. Then, once we have the proper transformation in hand, we can evaluate what happens to the fundamental objects under it, and make our inference to the spacetime structure undergirding the theory.

Consider a different case in which a theory's geometric formulation brings to light the underlying spacetime structure: Newtonian mechanics. Newton's first law says that every object continues in a state of rest or uniform motion unless acted on by a net external force. Geometrically, it says that every object continues on a straight spacetime trajectory unless acted on by a net force. Stated this way, Newton's law clearly requires that there be a genuine distinction between the trajectories of objects at rest or in uniform motion, and the trajectories of objects that are accelerating. If there were no objective, frame-independent distinction between accelerated and unaccelerated trajectories, then the law would not make any sense. Hence a world governed by Newtonian mechanics must possess the spacetime structure needed to ground that distinction. Such a world requires sufficient structure to distinguish, in an objective, frame-independent way, straight from curved paths through spacetime (cf. Maudlin 1994, Chapter 2).

Here is our guiding inference, put in these terms. Formulate our best fundamental theory in an invariant, geometric, coordinate-independent way. Then see what spacetime structure we need to assume in order to do this. The structure we need is the structure we infer the spacetime to have. Initially, make use of any structure we like in formulating a theory, including a temporal orientation structure. Just be sure to go back and check whether we really needed it.

How do we figure out whether we needed some bit of structure, such as a temporal orientation structure? Trot out the relevant transformation—in this case, our time reversal transformation. Apply it to the theory and see what happens. If the theory looks different after undergoing the transformation, then it seems to require or presuppose such a structure. If the theory looks the same after the transformation, then it does not seem to presuppose this structure; it turns out we did not need it to formulate the theory after all, and we correspondingly infer that there is none. The proper time reversal transformation should thus tell us whether we need a temporal orientation in order to formulate a theory in an invariant way; if we do, then we have reason to infer that there is one.[10]

Let us revisit the standard transformation, put in these terms. This

10. Keep in mind the distinction between the *coordinate*-invariant formulation of a theory, i.e., its invariance under allowable coordinate changes, and a theory's invariance under *time reversal*.

https://doi.org/10.1086/590199 Published online by Cambridge University Press

means updating our discussion to spacetime. For the time being, stick with the assumption of a flat spacetime.

The mirroring action is now this: flip the fundamental contents of spacetime across a time slice. Notice this is still captured by the idea of a movie playing backward, since the mirroring of objects across a time slice is the restacking of time slices in reverse temporal order. This is the spacetime equivalent of playing out the time-reversed sequence of states.

To apply this transformation, choose a frame within which to do the flipping, and choose a time slice within that frame to be the temporal origin. This time slice is our 'mirror'. Then shift the fundamental objects, or states, to their temporal mirror-image locations across the time slice mirror. For example, shift particle and scalar field value locations from (t, x, y, z) to $(-t, x, y, z)$. Passively, flip the time axis across the time slice mirror, changing these coordinate locations in the same way.

It is worth taking a moment to see how this transformation acts on other kinds of objects in spacetime. For now, focus on the active version. Since the standard transformation, understood actively, shifts particles to their temporal mirror image locations, this means that it will flip particle worldlines, taking their mirror images across the time slice mirror. This is the spacetime analog of particle locations getting mapped to their temporal mirror image locations in classical three-dimensional space.

It is a bit trickier to see what happens to 4-velocities. The 4-velocity of a particle at a spacetime point is the tangent to the worldline at that point.[11] In addition to flipping particle worldlines, the standard transformation also reflects their 4-velocities across the time slice mirror. The result is that their spatial components flip sign and their temporal components remain the same.

One might wonder why. One might think that the temporal component ought to flip. We are, after all, applying a *time* reversal transformation. Though the books say that the spatial components invert, they do not give much justification for this other than noting the desired result for 3-velocities: their spatial directions flip sign, just as they should in a backward-running movie.

Malament provides a justification for this by pointing to an important feature of 4-velocities. (See also Arntzenius 2005.) The worldline of a particle at a time really has *two* tangents associated with it, one pointing in each temporal direction. This means that there is no unique 4-velocity associated with a given point on a worldline. How then do we get the 4-velocity of a particle at a time? We need a temporal orientation on the

11. The 4-velocity is the geometric object representing instantaneous velocity in a four-dimensional spacetime. If a worldline is parameterized by the proper time, then the 4-velocity at a point is a unit vector tangent to the worldine at the point.

https://doi.org/10.1086/590199 Published online by Cambridge University Press

spacetime manifold. A temporal orientation picks out, at each spacetime point, which direction is to the future of that point and which is to the past. Given a temporal orientation, we can then pick out the unique future-directed 4-velocity at any point on a worldline—future-directed according to the background temporal orientation.

I say below why this gives rise to the standard transformation properties of 4-velocities. The important lesson for us is this. Whereas tangents to worldlines are defined without reference to a temporal orientation, 4-velocities are only defined relative to a temporal orientation and a worldline. Absent a temporal orientation, there simply is no unique 4-velocity associated with a point on a worldline. In this sense, 4-velocities are not truly fundamental objects. This will be important when we get to Malament's view.

(Why do the spatial components of 4-velocities invert? Assume a background temporal orientation, and so a unique 4-velocity associated with any point on a worldline. The standard active transformation mirrors worldlines across the time slice mirror, with the result that a 4-velocity's spatial components get the opposite sign. The reason is that the 4-velocity must remain tangent to the time-reversed worldline [at the corresponding time-reversed point], and the time-reversed worldline is the mirror image of the original worldline. The temporal component will not change because the time-reversed 4-velocity must still point to the future according to the fixed background temporal orientation. Since 4-velocities are *defined* relative to a temporal orientation and a worldline, and since the temporal orientation remains fixed while worldlines get flipped, this is what should happen to 4-velocities. Note that this yields the usual transformation properties of 3-velocities: since a 3-velocity comprises the three spatial components of the 4-velocity [the 4-velocity projected onto the x-y-z hypersurface], to say that the 3-velocity inverts [spatial] direction under time reversal is just to say that the three spatial components of the 4-velocity invert. But now we see a justification for these transformation properties of 4-velocities, over and above their yielding expected results for time-reversed 3-velocities.)[12]

6. Malament's Transformation. Malament's time reversal transformation is simple. Here it is in a nutshell: invert the temporal orientation.

That is it. Do not flip the contents of spacetime across a time slice. Do

12. Since 4-forces and 4-accelerations are defined independently of a temporal orientation, under standard time reversal the temporal components of these objects invert and their spatial components do not. It is now easy to see how the electromagnetic field behaves under standard time reversal, as Malament shows. See also Arntzenius 2005.

https://doi.org/10.1086/590199 Published online by Cambridge University Press

not shift particles and field values to their temporal mirror image locations. Do not invert particle worldlines. *Only* flip the temporal orientation vector field.

In the next section, I give reasons for preferring this transformation to the standard one. Here, I outline the way in which it is quite natural, details aside.

It is natural for a very simple reason. What is a time reversal transformation? Just a flipping of the direction of time! That is all there is to a transformation that changes how things are with respect to time: change the direction of time itself. In Malament's words, "The time reversal operation is naturally understood as one taking fields [and other fundamental objects] on [the spacetime manifold] M as determined relative to one temporal orientation to corresponding fields on M as determined relative to the other" (2004, 306).

This yields the following procedure to figure out whether a theory requires a temporal orientation. Take the theory in its coordinate-invariant formulation. Flip the temporal orientation (this is the step of applying the time reversal transformation). See whether the theory remains the same afterward. In other words, invert the direction of time, and see whether things still behave in accord with the theory. If they do, then the theory does not pick out one temporal direction as opposed to the other; it says the same thing regardless of the direction of time. In that case, it does not presuppose that there is a real distinction between the two directions. If things do look different afterward, then the theory does require the distinction, for it says different things depending on the direction of time. In this case, the theory does presuppose a real distinction between the two temporal directions, in only one of which things can lawfully evolve. And a temporal orientation is the spacetime structure that supports such a distinction.

As a word of caution, let me emphasize that more needs to be done to in order show that the procedure will work as a general rule. It may not always be clear what role the temporal orientation is playing in a given theory's invariant formulation. Malament works this out for the particular case of classical electromagnetism, but it remains to be seen whether the procedure will work for any theory we take to be time reversal invariant.[13]

13. In classical electromagnetism, the key is Malament's realization that we can characterize the electromagnetic field without reference to a temporal orientation, as a linear map from tangent lines to 4-forces (or 4-accelerations). The standard characterization of the electromagnetic field in terms of the tensor field F_{ab} does make reference to a temporal orientation; this, too, falls out of Malament's formulation, once we introduce a background temporal orientation.

https://doi.org/10.1086/590199 Published online by Cambridge University Press

7. Two Views. The standard time reversal transformation mirrors the fundamental objects in spacetime across a time, keeping the background temporal orientation (if there is one) fixed. A theory is time reversal invariant if this does not alter the theory: a solution always gets mapped to a solution. Malament's transformation leaves the fundamental objects in spacetime alone and flips the temporal orientation. A theory is time reversal invariant if a solution always remains a solution after we invert the temporal orientation.

In this section, I compare Malament's view with the standard view on six points, concluding that Malament's transformation is the one we should use in deciding whether a given theory is time reversal invariant. Before doing so, let me first note some features of it that may seem counterintuitive.

Malament's approach seems to require that we treat the temporal orientation as a real physical field. As Malament discusses it, the temporal orientation takes on a life of its own as a fundamental vector field on the spacetime manifold. Indeed, there is no such thing as a 4-velocity, let alone a time reversal transformation, without it.

One might balk at this understanding of the temporal orientation. It is certainly an odd-sounding sort of field. Should we posit a physical field corresponding to *any* noninvariance of a theory? If our theories are noninvariant under Galilean transformations, should we infer an 'absolute rest vector field' on the spacetime? An even odder-sounding field! Further, for a temporally symmetric theory like Newtonian mechanics, we want to say that there simply *is* no temporal orientation field. How then can we define the time reversal invariance of a theory in terms of this (nonexistent) field?

My own view is that this is the right way to understand these coordinate-independent objects in general, and the temporal orientation in particular. We need these abstract-seeming objects in our (best formulations of) fundamental physics, and I see no reason to distinguish between a temporal orientation, on the one hand, and an electromagnetic field—or, for that matter, a metric or other spacetime structure—on the other, as far as their being physically real goes (North 2008). Nevertheless, one does not have to adopt this understanding of the temporal orientation in order to accept Malament's basic notion of time reversal. Just apply your preferred view of the abstract geometric objects used in physics, and the physical objects they represent, to this case. And when we apply our time reversal test, we do not have to (absurdly) posit the existence of a nonexistent field. Rather, we are considering what would happen to the theory if there were such a field. If the theory would be the same regardless, then we infer that there is no such field. If the theory would be different, then we infer that this is because there is such a field.

https://doi.org/10.1086/590199 Published online by Cambridge University Press

Second, on Malament's approach, there is no sense to be made of 4-velocities absent a temporal orientation. This might seem odd as well. But it is correct: the usual definition of a 4-velocity just does not allow for a unique such object without a temporal orientation.[14] And note that this does not mean that we are left without any well-defined 4-velocities in a nontemporally oriented spacetime. It is open to conventionally choose a temporal orientation if the spacetime does not come equipped with one. (This does mean that we cannot define 4-velocities in a nontemporally orientable spacetime, but such a spacetime is so strange anyway that this need not worry us here.)[15]

Third, Malament's view goes against what I presented as the intuitive understanding of time reversal. Malament's time reversal transformation keeps every fundamental object in spacetime at the same location. Nothing gets shifted at all: we only flip the temporal orientation vector field. This means that particle positions and worldlines remain put, instead of shifting to their temporal mirror-image locations as they would in a reverse-running movie; clockwise motions are not time-reversed into counterclockwise motions, for example.[16] In reply, I think we need only point to the naturalness of this sense of time reversal, discussed in the last section, and to the arguments coming up in this section. The backward-running movie idea is a relic left over from classical theories set in three-dimensional space. (Think of the backward-running movie as the inverted sequence of movie frames, i.e., the inverted history of objects' spatial locations at the same time.)

Now to compare the views on six points. Together, these points make the case for Malament's transformation over the standard view, though the last two alone should suffice.

7.1. Justification. One problem with the standard view is its lack of a

14. We could define a 4-velocity, or alternatively a worldline, in a nonstandard way, as an intrinsically temporally directed object. I am open to this idea, though it is both intuitively odd and highly nonstandard. See Arntzenius and Greaves (2007) for discussion.

15. There are spacetimes of general relativity that are not globally temporally orientable. Our own spacetime locally has a Lorentz signature metric: locally, it is temporally orientable. This does not guarantee global temporal orientability. But as John Earman (2002, 257) argues, so long as we have good reason to think the fundamental laws hold in all regions of our universe, we have good reason to infer that local temporal orientability yields global temporal orientability.

16. Malament suggests that this falls out of his view given a reference frame (in his sense), though I am not sure that this preserves the initial idea. It seems to me that one of the main benefits of Malament's view (the applicability to arbitrary curved spacetimes) indicates that we will have to give up the classical movie-playing picture.

https://doi.org/10.1086/590199 Published online by Cambridge University Press

clear motivation. Physics books typically do not offer much justification other than to point to desired results for classical theories, as we saw above for their view on 4-velocities. Malament shows that there is a justification available, once we see that 4-velocities are defined relative to a temporal orientation. But this raises a question: Why not invert the temporal orientation as well? The temporal orientation can be thought of as a fundamental object, just as worldlines are. Then why does it not get flipped across the time slice mirror too?[17] The books do not say anything about this. They just assume that the temporal orientation remains fixed. (They are not explicit about this; Malament brings the assumption to light.)

Malament, though, has a clear motivation for his view. To figure out whether a theory requires a temporal orientation, compare the theory with what it says after flipping the temporal orientation. This understanding of time reversal yields a clear justification for objects' transformation properties. Take 4-velocities. Their transformation properties will differ from those on the standard view: since 4-velocities are defined relative to a temporal orientation and a worldline, when we invert the temporal orientation, keeping everything else (worldlines included) fixed, *all* the components will flip, temporal and spatial. More generally, when we flip the temporal orientation—which is all that this transformation does, remember—all the components of objects defined relative to a temporal orientation will flip sign; the components of objects not so defined will not. Thus, fundamental objects (other than the temporal orientation itself) are left alone by the transformation; those defined relative to a temporal orientation completely invert. Worldlines are invariant, as are their tangents, whereas 4-velocities completely invert. Indeed, only objects defined relative to a temporal orientation change at all under this transformation.[18]

So the reason that 4-velocities completely invert on Malament's view is the same as the reason that only their spatial components invert on the standard view: 4-velocities are defined relative to a temporal orientation and a worldline. The difference in transformation properties comes from

17. This might suggest a third option: flip worldlines and the temporal orientation. But this transformation will always give us our original states right back. The time-reversed solution is just the original solution, looked at from a different perspective.

18. 4-forces and 4-accelerations will not change any components. As a result, classical electromagnetism is time reversal invariant on this transformation, as it is on the standard one. Malament locates a conception of the electromagnetic field that differs from the standard Maxwell-Faraday tensor (as a map from tangents to 4-forces), which is independent of a temporal orientation. Given this and the above characterization of 4-velocities, it is easy to see how the standard electromagnetic field tensor transforms under time reversal.

https://doi.org/10.1086/590199 Published online by Cambridge University Press

the difference in the basic action of the time reversal operation. Malament's transformation flips the temporal orientation and leaves worldlines intact; the standard view flips worldlines and leaves the temporal orientation intact. Yet Malament has a clear justification for the transformation properties of the different objects in spacetime, based on how these objects are defined and what the time reversal operation does. The standard view has a justification for the transformation properties of some objects, like 4-velocities, but this only reveals the lack of a justification for other aspects of the view, such as its leaving the background temporal orientation fixed.

7.2. Active versus Passive Time Reversal. In a classical space, there is a passive time reversal transformation corresponding to the standard active one, as we saw above. This is less straightforward in a four-dimensional spacetime. In fact, now the standard view seems to lose the correspondence between active and passive time reversal entirely. Recall that the standard active transformation mirrors the material contents of spacetime across a time slice. What is the corresponding passive transformation?

The passive transformation should just be a change in time coordinate from t to $-t$. Yet this would yield different time-reversed 4-velocities, inverting their temporal and not their spatial components. Take any 4-velocity V_μ, say with original components (1, 2, 3, 4) in a given coordinate system. Under standard active time reversal, the components become (1, −2, −3, −4), whereas under this passive transformation, they become (−1, 2, 3, 4).

This is weird. Usually there is a passive transformation corresponding to any given active one. And, intuitively, it should not matter whether we change the coordinate-dependent representations of invariant objects by shifting the objects themselves or by moving the coordinate system around them. Not only is this intuitively odd, but it seems to undercut our inference to a direction of time. Since the active and passive transformations have different effects on objects in spacetime, whether a theory is deemed time reversal invariant or not could depend on whether we use the active or passive transformation as our test. Our inference to a temporal orientation would then similarly depend on this.

Since the books assert that 4-velocities' spatial components flip sign under time reversal, it seems the standard transformation must be understood actively. The passive transformation simply will not yield the result the books claim for time-reversed 4-velocities. We could try to construct a passive transformation to yield the same results as the active one. But a passive transformation is *defined* as a change in coordinate-dependent representations of coordinate-independent objects *due to a change in coordinate system*. To define a passive transformation in terms of an active one is to give up on the distinction between active and passive

https://doi.org/10.1086/590199 Published online by Cambridge University Press

transformations altogether.[19] Unlike in classical space, then, here we cannot understand the standard transformation (passively) as an inversion of the time axis. Instead, both the chosen coordinate system and the background temporal orientation must remain intact.

Here is another way of defending the standard view of 4-velocities, and by means of what is perhaps a passive version of the standard transformation.[20] Consider a transformation that flips the time coordinate, not the temporal orientation or worldlines, and also says that the proper time goes with the time coordinate. Once we flip the temporal coordinate, the parameterization of a worldline by the proper time will also flip. The result (assuming that a 4-velocity is defined relative to a parameterization by proper time) is that a 4-velocity's spatial components flip sign, and the temporal component does not, just as the standard view says. We can regard this as a kind of passive transformation corresponding to the standard active one, in this case changing not just t to $-t$ but also the parameterization. (I say "kind of" since this does not just change the coordinate system, and I am not sure whether this should count as active or passive.) On this view, too, 4-velocities are not fundamental objects: they are defined relative to a parameterization. This is another way of defending the standard view, and with something that looks like a passive transformation—it does not shift around fundamental objects in spacetime—albeit not the usual one that only changes t to $-t$. Nonetheless, this will also face the two key problems below.

Malament's time reversal is also an active transformation: invert the temporal orientation, understood as a real vector field, represented by an abstract geometric object, on the spacetime manifold. Here, too, it seems there simply is no corresponding passive transformation.[21] No passive transformation will get the transformation properties for objects in spacetime that Malament's time reversal does. Think of inverting the time axis. Whereas Malament's transformation flips all the components of 4-velocities, this passive transformation flips only the temporal component. In general, geometric objects transform differently under Malament's time reversal depending on whether they are defined relative to a temporal orientation

19. Earman (2002) defines time reversal as a change in coordinate system plus a flipping of the temporal orientation. This too mixes the notion of passive with active time reversal in such a way as to lose the distinction.

20. This was pointed out to me by Hartry Field (in conversation) and is advocated by Earman (2002, Section 3).

21. But see Leeds 2006 for an argument that Malament's transformation is first and foremost a passive transformation, with an active version.

https://doi.org/10.1086/590199 Published online by Cambridge University Press

or not, and there is no passive transformation, no mere change in coordinate system, that will get this result.[22]

Thus, both views face the problem of losing the correspondence between active and passive time reversal. On this, the views are tied. I am inclined to think this is not such a problem for either view. Once we start investigating theories in their coordinate-invariant formulations, it seems their symmetry properties ought to be linked to transformations that do their job independently of an arbitrarily chosen coordinate system. Since the theories themselves are written in a way that is invariant under coordinate transformations, it is reasonable to expect that the proper time reversal (or other) transformation should be too.

7.3. Time Translation Invariance. Standard time reversal presupposes that theories are time translation invariant.[23] The standard transformation mirrors objects across a chosen temporal origin. If a theory is not time translation invariant, then whether it is deemed time reversal invariant will differ depending on the choice of temporal origin. Indeed, the standard picture of time reversal is a bit unclear if a theory did depend on the time. During what time sequence should we compare an allowable process with its time reverse?

Although the theories we take seriously for our world all appear to be time translation invariant, it is a benefit of Malament's view that it does not presuppose this. Since Malament's transformation does not require any choice of temporal origin, it will deem a theory time reversal invariant, or not, regardless of whether it possesses this additional symmetry.

7.4. Indeterministic Theories. Malament's test for time reversal can be easily adapted to indeterministic theories. Formulate the theory in a coordinate-independent way. Take any (conditional) probability the theory assigns to a history of states. Flip the temporal orientation. See whether the probabilities remain the same. If they do, then the theory is time reversal invariant, and we did not need a temporal orientation to formulate it in the first place.

Although the standard transformation can be made to work for in-

22. Keep in mind that flipping the temporal orientation is different from inverting the time coordinate axis. (1) The temporal orientation is a fundamental coordinate-independent object, represented by a vector field on the spacetime; the time axis is a coordinate axis we lay down on the spacetime in order to represent invariant objects. (2) Inverting the time axis has a different effect on 4-velocities from inverting the temporal orientation. The standard view does keep both the time axis and the temporal orientation fixed, but not because these must go together. It is because the view inverts worldlines and not the temporal orientation.

23. I thank Matt Kotzen for pointing this out to me.

https://doi.org/10.1086/590199 Published online by Cambridge University Press

deterministic theories, the story is more involved. I refer the reader to Arntzenius 1997 for that story.

7.5. Curved Spacetimes. Malament's transformation applies to background spacetimes of arbitrary curvature. This means it is suitable for the spacetimes of general relativity. The only requirement is that the spacetime be temporally orientable.

The standard transformation will not generalize in this way. It flips fundamental objects across a time slice, (actively) shifting them to spacetime locations at an equal temporal distance away from the time slice mirror. This requires a choice of reference frame in which to do the flipping and a choice of time slice mirror within that frame. The mirroring action must shift objects' locations to an equal temporal distance away from this time slice, while retaining relative distances along the mirror-image, time-reversed worldlines.

All of this requires that we can measure an equal distance in space from a given time. And it is assumed that we can do this along a straight line. But the notions of distances along a straight line and of equal distances from a time straightforwardly make sense only in a flat spacetime. The standard picture requires a background spacetime with enough time-reflection symmetries that we can mirror objects across a time slice in this way. This suggests that it presupposes a spacetime with the full structure of a Minkowski spacetime.

We could try using distances along a spacetime's geodesics instead of distances along straight lines; geodesics are the curved-spacetime analogs of straight lines in a flat spacetime. But this will not work as a general rule. Even if the geodesics are orthogonal to our time slice mirror (the spacelike hypersurface we choose as our temporal origin), they may not be orthogonal to a later time slice. In general, geodesics have all sorts of strange properties that do not fit nicely with the standard idea of time reversal as a mirroring across a time. Geodesics can cross, curve, not remain parallel, they can intersect each other more than once. This means, for instance, that two distinct spacetime points could be mapped to the same point under time reversal, and we would not want that to happen: two distinct points that are simultaneous in a frame should remain simultaneous after time reversal, not get mapped to the same spacetime location. In short, it is crucial to the standard view that two straight parallel lines do not intersect, and taking temporal mirror image locations along geodesics will not guarantee this, for geodesics can cross.

We could construct a later time slice by following the spacetime's geodesics, making sure that any time slice along the way remains orthogonal to them. This could also yield time-reversed locations at equal temporal distances away from the time slice mirror as the original locations are.

https://doi.org/10.1086/590199 Published online by Cambridge University Press

Typically, though, the hypersurface at the later time will itself be curved and crossed over itself, since the geodesics we use to construct it can do this. But then the hypersurface will no longer reasonably be a time slice! There simply is no guarantee of getting flat hypersurfaces over time—genuine *time slices*—by constructing them in this way. (We could find time slices by means of Killing vector fields orthogonal to the geodesics, but that is a special case. We cannot in general construct hypersurfaces through an arbitrary curved spacetime so that the geodesics in any direction will be orthogonal to the time slices.)[24]

The standard transformation therefore presupposes a flat background spacetime (modulo the above). Moreover, it is defined relative to a frame of reference and a time slice within it. Objects then have their transformation properties relative to that frame; the spatial components of 4-velocities flip sign within the frame in which we choose to do the flipping, for example. This is to be expected from a transformation that tries as much as possible to preserve the classical picture of time reversal. Consider two objects moving with constant relative velocity in a classical spacetime: the time-reversed velocities clearly depend on the frame in which we choose to carry out the time reversal.

Malament's transformation generalizes to the (temporally orientable) arbitrary curved spacetimes of general relativity, while yielding the same results as the standard view for spacetimes that have the above symmetries. Further, Malament's is an entirely frame-independent notion of time reversal. Objects completely invert, in any frame, if they are defined relative to a temporal orientation, or remain invariant, in any frame, if they are not so defined. We are always free to arbitrarily choose a frame in order to employ the standard transformation. In discussing the view above, we assumed that such a choice had been made. But it is certainly cleaner if we do not have to do this.

24. As Malament pointed out to me (in correspondence), there are non-Minkowski spacetimes that allow for a limited sense of standard time reversal, namely those admitting one well-defined family of time-reflection maps, by means of a hypersurface-orthogonal timelike Killing vector field. Hypersurface-orthogonality means that (at least locally) we can foliate the spacetime with a family of spacelike hypersurfaces everywhere orthogonal to the Killing vector field. Even so, only in a Minkowski spacetime will there be many hypersurface-orthogonal timelike Killing fields, allowing for many families of time reflection maps, which we can apply to arbitrary motions. (Malament offers this precisification as a proposition: Let (M, g_{ab}) be a (not necessarily temporally orientable) spacetime. Suppose that for all points p in M, and all timelike vectors ξ_a at p, ξ_a can be extended to a hypersurface-orthogonal Killing field on M. Further suppose that (M, g_{ab}) is geodesically complete. Then (M, g_{ab}) is isometric to Minkowski spacetime: locally, the timelike Killing fields means the spacetime metric is flat, with vanishing Riemann curvature.)

https://doi.org/10.1086/590199 Published online by Cambridge University Press

7.6. The Inference to a Direction of Time. Malament's transformation yields a better test for a temporal orientation, the very thing we want from our time reversal transformation. There are two ways to see this.

The first is its naturalness. Malament's test goes like this. Take our best fundamental theory and take a solution to it. Invert the temporal orientation and see whether it remains a solution. If it always does, then the theory did not require the temporal orientation to begin with. It says the same thing whether time goes in one way or the other. If a solution does become a nonsolution, then this is indication that the temporal orientation matters to the theory. After all, this is the only thing that changed under the transformation.[25]

Second, Malament's time reversal transformation, unlike the standard one, does not invert spacetime handedness. To see this, think of a spatiotemporally handed object, say in a three-dimensional spacetime.[26] Suppose the object comprises three arrows, one pointing up along the temporal axis (call this component arrow **t**), the other two in the *x*-*y* plane (call these **x** and **y**). This is a triple of linearly independent vectors, one timelike, the other two spacelike. Suppose this is a fundamental object, not defined relative to any other fundamental objects in the spacetime.

Now apply the standard time reversal transformation. This transformation (actively) mirrors fundamental objects across a space of one dimension fewer than the space the objects are in (and which is orthogonal to the time axis). Suppose we choose the *x*-*y* plane to be the time slice mirror, and then mirror the object across this plane.

What happens when we do this? We invert the object's spacetime handedness! (Note that it is the *spacetime* handedness, not spatial handedness, that is inverted.) After time reversal, the two spacelike arrows in the *x*-*y* plane remain where they are, and the timelike arrow points in the opposite temporal direction. Before time reversal, when we rotated **x** into **y**, we got **t** pointing up: this was a right-handed object (right-spatiotemporal-handed). After time reversal, we can no longer rotate **x** into **y** and get **t** pointing up: this is a *left*-handed object.

Any mirroring across a hypersurface of one dimension fewer than the

25. Callender (2000) gives a spatial analogy. Imagine a cube-shaped fundamental particle that emits photons from one of its faces during an experiment. Repeat the experiment, with the only difference being that the cube is spatially mirror-reflected. If the photons are now emitted in the same spatial direction rather than the mirror-image one, it seems reasonable to conclude that space itself has a preferred direction, 'pulling' the photons that way.

26. A handed object cannot be mapped to its mirror image via ordinary rotations and translations within the space (a right-handed glove in a three-dimensional space, the letter *L* written on a chalkboard). In any space, there are only two possible handednesses: handed objects either have the same or opposite handedness from one another.

https://doi.org/10.1086/590199 Published online by Cambridge University Press

space—which is what the standard time reversal operation *is*—will wind up inverting spacetime handedness as well. Actively: flip the fundamental objects across a time, and we change the spacetime handedness of any objects that are spacetime handed. Passively: flip the time axis about the origin, and we change the spacetime handedness of the coordinate system.[27]

This means that if a theory is not invariant under standard time reversal, we can reasonably infer that it requires *some* asymmetric spacetime structure allowing things to evolve in only one of two ways. But this could be a spacetime-handedness structure rather than an asymmetric temporal structure. If a theory changes under this transformation, the culprit could be *either* a direction of time *or* a spacetime-handedness structure (a spacetime-handedness orientation field, say). So this is a bad test, if what we are ultimately after is evidence of a temporal orientation.

Under Malament's transformation, no fundamental spacetime object, including the one above, will undergo any change that could fiddle with its spacetime handedness (if it has such a handedness). True, after time reversal, part of the object will point to the past rather than the future. But its overall spacetime-handedness will not change.[28] The reason is that this transformation leaves *everything*—every fundamental object in spacetime, other than the direction of time itself—intact.[29]

Malament's time reversal transformation yields a good test for a direction of time precisely because it *only* flips the temporal orientation. If things look different afterward, then the only possible culprit is the temporal orientation. Since the direction of time is the only thing altered in the time-reversed scenario, it is natural to conclude that it plays a real role in the theory. If this is our best, most fundamental theory, then this gives us reason to infer that the spacetime itself is temporally oriented.

8. Conclusion. We should use Malament's time reversal transformation to test whether a given theory is time reversal invariant, and to infer

27. If the active and passive transformations corresponded, then flipping the spacetime handedness of coordinates would be equivalent to flipping the objects' spacetime handedness.

28. The spacetime handedness of nonfundamental objects might change, but that will not affect the time reversal invariance of our fundamental theories.

29. You might think that the spacetime handedness must change once the temporal orientation flips. Keep in mind that a handedness is just an orientation or an ordering, where permutations change the handedness (see note 26). Suppose an ordered triple $(1, 2, 3)$ is right-handed according to the background spatiotemporal orientation. After inverting the temporal orientation, this object should have the same ordering, that is, the same spacetime handedness. Malament (2004, Sections 5–6) discusses this from the four-dimensional geometric perspective, which makes this clear. See Arntzenius and Greaves 2007 for more discussion.

https://doi.org/10.1086/590199 Published online by Cambridge University Press

whether the theory's spacetime is temporally oriented. Unlike the standard transformation, Malament's time reversal does not mix in considerations of spacetime handedness or other symmetries, and it generalizes to the arbitrary curved spacetimes of general relativity. The only drawback is that it fails to respect our initial intuition about time reversal. But this is insufficient reason for clinging to the standard view, especially when, on further reflection, it is Malament's transformation that captures what we had in mind all along: to see how a theory fares when oriented with respect to one direction of time as opposed to the other.

REFERENCES

Albert, David Z (2000), *Time and Chance*. Cambridge, MA: Harvard University Press.

Arntzenius, Frank (1995), "Indeterminism and the Direction of Time", *Topoi* 14: 67–81.

——— (1997), "Mirrors and the Direction of Time", *Philosophy of Science* (Proceedings) 64: S213–S222.

——— (2004), "Time Reversal Operations, Representations of the Lorentz Group, and the Direction of Time", *Studies in History and Philosophy of Modern Physics* 35: 31–43.

——— (2005), "Are Anti-particles Just Particles Traveling Back in Time?". Unpublished manuscript.

Arntzenius, Frank, and Hilary Greaves (2007), "Time Reversal in Classical Electromagnetism". Unpublished manuscript.

Callender, Craig (2000), "Is Time 'Handed' in a Quantum World?", *Proceedings of the Aristotelian Society*: 247–269.

Earman, John (2002), "What Time Reversal Invariance Is and Why It Matters", *International Studies in the Philosophy of Science* 16: 245–264.

Horwich, Paul (1987), *Asymmetries in Time: Problems in the Philosophy of Science*. Cambridge, MA: MIT Press.

Leeds, Stephen (2006), "Discussion: Malament on Time Reversal", *Philosophy of Science* 73: 448–458.

Malament, David (2004), "On the Time Reversal Invariance of Classical Electromagnetic Theory", *Studies in History and Philosophy of Modern Physics* 35: 295–315.

Maudlin, Tim (1994), *Quantum Non-locality and Relativity: Metaphysical Intimations of Modern Physics*. Oxford: Blackwell.

——— (2006), "Relativity", in Donald M. Borchert (ed.), *Encyclopedia of Philosophy*. Detroit: Thompson Gale.

——— (2007), "On the Passing of Time", in *The Metaphysics within Physics*. Oxford: Oxford University Press, 104–142.

North, Jill (2008), "The 'Structure' of Physics: A Case Study". Unpublished manuscript.

Price, Huw (1996), *Time's Arrow and Archimedes' Point: New Directions for the Physics of Time*. Oxford: Oxford University Press.

Savitt, Steven (1996), "The Direction of Time", *British Journal for the Philosophy of Science* 47: 347–370.

https://doi.org/10.1086/590199 Published online by Cambridge University Press